Checking if a PBS Job was Killed by the OOM Killer

If a PBS job runs out of memory and is killed by the kernel's out-of-memory (OOM) killer, the event is usually recorded in system log files. If the log files confirm that your job was killed by the OOM killer, you can get your job running by increasing your memory request.

Complete the following steps to check whether your job has been killed by the OOM killer:

1. Run the **tracejob** command to find out when your job ran, what rack numbers were used, and if the job exited with the **Exit_status=137** message. For example, to trace job 140001 that ran within the past three days, run:

   ```
   pfe21% ssh pbspl1
   pbspl1% tracejob -n 3 140001
   ```

   where **3** indicates that you want to trace a job that ran within the past 3 days and **140001** indicates the **job_id**. (**pbspl1** is the PBS server for Pleiades, Aitken, and Electra; for Endeavour, use **pbspl4**.)
2. In the **tracejob** output, find the job's rack numbers (such as r2, r3, ...), then run **grep** to find messages that were recorded in the messages file, which is stored in the leader nodes of those racks. For example, to view messages for rack r2:

   ```
   pfe21% grep abc.exe /net/r2lead/var/log/messages
   Apr 21 00:32:50 r2i2n7 kernel: abc.exe invoked oom-killer:
   gfp_mask=0x201d2, order=0, oomkilladj=-17
   ```

   Note: Often, an out-of-memory message will not be recorded in the messages file, but will be recorded in a consoles file named by each individual node. In that case, you'll need to **grep** the messages in the consoles file. For example, to look for **abc.exe** invoking the OOM killer on node r2i2n7:

   ```
   pfe21%  grep abc.exe /net/r2lead/var/log/consoles/r2i2n7
   abc.exe invoked oom-killer: gfp_mask=0x201d2, order=0, oomkilladj=0
   ```
3. Use an editor to view the files and look for the hourly time markers bracketing the period of time when the job ran out of memory. An hourly time marker looks like this:

   ```
   [-- MARK -- Thu Apr 21 00:00:00 2011]
   ```

Note: Sometimes a system process (such as **pbs_mom** or **ntpd**) is listed as invoking the OOM killer; this is also direct evidence that the node ran out of memory.

TIP: We provide a script, called **pbs_oom_check**, which does these steps automatically and also parses the **/var/log/messages** on all the nodes associated with the **jobid** to search for an instance of OOM killer. The script is available in **/u/scicon/tools/bin** and and works best when run on the host pbspl1.
If your job needs more memory, see How to Get More Memory for Your PBS Job for possible approaches. If you want to monitor the memory use of your job while it is running, you can use the tools listed in Memory Usage Overview.

Article ID: 221
Last updated: 13 May, 2021
Revision: 34
Running Jobs with PBS -> Optimizing/Troubleshooting -> Checking if a PBS Job was Killed by the OOM Killer
https://www.nas.nasa.gov/hecc/support/kb/entry/221/